
A Bayesian Approach to Predicting Disengaged Youth

David Kohn

Centre for Translational Data Science
University of Sydney
New South Wales 2006
david.kohn@sydney.edu.au

Sally Cripps

Centre for Translational Data Science
University of Sydney
New South Wales 2006

Nick Glozier

Brain Mind Centre
University of Sydney
New South Wales 2006

Hugh Durrant-Whyte

Centre for Translational Data Science
University of Sydney
New South Wales 2006

Abstract

This article presents a Bayesian approach for predicting and identifying the factors which most influence an individual's propensity to fall into the category of Not in Employment Education or Training (NEET). The approach partitions the covariates into two groups: those which have the potential to be changed as a result of an intervention strategy and those which must be controlled for. This partition allows us to develop models and identify important factors conditional on the control covariates, which is useful for clinicians and policy makers who wish to identify potential intervention strategies. Using the data obtained by O'Dea et al. (2014) we compare the results from this approach with the results from O'Dea et al. (2014) and with the results obtained using the Bayesian variable selection procedure of Lamnisos et al. (2009) when the covariates are not partitioned. We find that the relative importance of predictive factors varies greatly depending upon the control covariates. This has enormous implications when deciding on what interventions are most useful to prevent young people from being NEET.

1 Background

Often the number of factors available for prediction on a given individual, p , is larger than the number of individuals on whom we have NEET status measurements, n , making the identification of important factors statistically challenging. For example, inference in a frequentist procedure usually relies on the assumption of asymptotic normality of the sample estimates, however as $p \rightarrow n$ this assumption is unlikely to be true. Another related issue is that good predictive performance does not necessarily equate with the identification of causal factors; many different combinations of factors may be equally good in predicting whether or not an individual will end up in the NEET category. However, if issues such as NEET are to be addressed, policy makers need to know what modifiable factors are likely to be causal so that appropriate intervention strategies are used.

To address the joint issues of high dimensionality and causal inference we take a Bayesian approach. Specifically, we propose to reduce the dimensionality by using a spike and slab prior over the regression coefficients (see, for example, Lamnisos et al., 2009). This regularization may result in biased estimates of the regression coefficients as per Chernozhukov et al. (2016). To address this issue we develop a series of conditional models by dividing the covariates into two groups, those which have the potential to be changed by an intervention, for example an individual's clinically assessed depression score, and those which do not, for example an individual's age or sex. These

conditional models are typically based on very small sample sizes, with $p > n$, making variable selection difficult but important.

2 Model

Suppose we have observations on n individuals' NEET status, $\mathbf{y} = (y_1, \dots, y_n)$, where $y_i = 1$ if an individual i is classified as NEET and $y_i = 0$ otherwise and corresponding measurements on covariates for each of these individuals. We denote the potentially causal factors by X and the other factors, which we refer to as control factors, by W , with $X = (1, \mathbf{x}'_1, \dots, \mathbf{x}'_n)$, where $\mathbf{x}_i = (x_{i1}, \dots, x_{iP_x})$, and x_{ki} the measurement of the k^{th} covariate on individual i , for $i = 1, \dots, n$, and $k = 1, \dots, P_x$, and P_x is the number of covariates which are potentially modifiable. We model the dependence between y , conditional on the control variables, w , and x using a generalized linear model (GLM),

$$\Pr(y_i = 1 | \mathbf{w}_i, \mathbf{x}_i) = g(\mathbf{x}_i \beta(\mathbf{w}_i)), \quad (1)$$

where g is some link function, and the notation $\beta(w)$ means that the $(P_x + 1) \times 1$ vector of regression coefficients is parameterized to depend upon the control variables W . This paper uses the standard normal CDF as the link function, so that $g(X\beta(w)) = \Phi(X\beta(w))$, where $\Phi(X\beta(w)) = \Pr(z < X\beta(w))$ and $z \sim N(0, 1)$. The choice of Φ as the link function allows us to employ the data augmentation method of Albert and Chib (1993) to estimate the regression coefficients and perform variable selection.

To obtain the likelihood function we divide the data into $S \leq n$ non-overlapping partitions, $W = (\mathbf{w}_1, \dots, \mathbf{w}_S)$, where $\mathbf{w}_s = (1, w_{s1}, \dots, w_{sP_w})$ represents a $(P_w + 1) \times 1$ vector of unique values of the control factors, for $s = 1, \dots, S$, with P_w the number of control covariates. Let $n_s \geq 1$, be the number of observations in each partition $s = 1, \dots, S$ and define I_s to be the set of indices for observations corresponding to partition s for $s = 1, \dots, S$. Then, the likelihood function is

$$p(\mathbf{y} | W, X, \beta) = \prod_{s=1}^S \prod_{i \in I_s} \Phi(\mathbf{x}_i \beta(\mathbf{w}_s))^{y_i} (1 - \Phi(\mathbf{x}_i \beta(\mathbf{w}_s)))^{(1-y_i)}. \quad (2)$$

To fully specify the model we place priors on those parameters needed to evaluate the likelihood, namely the regression coefficients. We wish to place a prior on these regression coefficients to allow for the possibility that some causal factors on which we have measurements do not have an impact on the future NEET status of an individual and also to allow this possibility to depend upon the control factors.

Specifically, for each partition s , we introduce an indicator vector $\gamma_s = (\gamma_{1s}, \dots, \gamma_{P_x, s})$, where $\gamma_{ks} = 1$ if causal factor k is in the model for partition s and $\gamma_{ks} = 0$ otherwise. To write the prior for the regression coefficients, we define the set $A_{1s} = \{k : \gamma_{ks} = 1\}$ to be the set of indices corresponding to those causal factors which are included in the model for partition s and define A_{0s} similarly.

Finally, we define $\Pr(\gamma_{ks} = 1 | \mathbf{w}_s) = \pi_k(\mathbf{w}_s)$ to be the prior probability that the k^{th} causal factor is in the model for partition s . This is parameterized to depend on the control covariates \mathbf{w}_s . We model the dependence between γ_k and \mathbf{w}_s as a probit regression, so

$$\pi_k(\mathbf{w}_s) = \Phi(\mathbf{f}_k(\mathbf{w}_s)), \quad (3)$$

where f is some, possibly non-linear, function. For now we take f as linear so that $\mathbf{f}_k(\mathbf{w}_s) = \mathbf{w}_s \alpha_k$, where $\alpha_k = (\alpha_{k0}, \dots, \alpha_{kP_w})$, is the $(P_w + 1) \times 1$ vector of regression coefficients for (3).

3 Data Description

We use data from the 'Transitions Study' as used in O'Dea et al. (2014) and detailed in Purcell et al. (2015). The study was conducted at two time periods, 'baseline' and 'followup' and implemented at four mental health service centres in Australia. The 377 participants in the sample are aged 16-25 and were subject to a variety of demographic questions and clinical and psychological assessments. Our target variable is NEET status, defined as not being in employment, education or training in the past month, measured at both baseline and follow up periods. Our other variables are measured only at the baseline period and include QIDS, which assesses the presence of major diagnostic symptoms

of depression; WHO-ASSIST, which assesses risky use of tobacco, alcohol and cannabis; GAD, which measures the symptoms of generalized anxiety disorder; WHODAS, which is a self rated examination of perceived functioning in daily life; and the demographic factors age and sex. To make interpretation easier we refer to QIDS, GAD and WHODAS as depression, anxiety and functioning respectively.

4 Results

In this analysis we consider two settings for the response variable, (i) an individual's NEET status at baseline and (ii) an individual's NEET status at follow up, and two settings for the covariates (i) using depression, tobacco, alcohol, cannabis, anxiety, functioning, age and sex as potential predictive factors (ii) using only the "modifiable" factors depression, tobacco, alcohol, cannabis, anxiety, functioning as potential predictors. To compare our method with other techniques we first analyse the data using commonly used Bayesian and frequentist techniques.

For the Bayesian variable selection analysis we follow Lamnisos et al. (2009) and use a g -prior prior over the regression coefficients. The results of this analysis appear in Table 1. Panel (a) shows the results using baseline NEET as the target variable while panel (b) uses followup NEET as the target variable. The quantity $\hat{\beta}$ is estimate of the posterior mean of β conditional on $\beta \neq 0$, i.e. $E(\beta|y, \beta \neq 0)$. The MPP label refers to the marginal posterior probability of inclusion for a variable.

Table 1 shows that the covariates generally have higher MPPs for the baseline NEET analysis relative to the followup NEET analysis, indicating that the concurrent measurements are correlated with NEET status but are less relevant as temporal, causal factors. Sex and age have high MPPs in the baseline model, motivating our conditioning on these variables as control factors so we can get a more finely gridded estimate of the probability of inclusion for the modifiable factors. The accuracy of the models, as given by the area under the ROC curve, (AUC), estimated by 5 fold crossvalidation and seen in the bottom row of table 1, is similar. 5-fold crossvalidation on the average model for both baseline and followup results in relatively similar ROC area-under-curve (AUC) for all models, as seen in the bottom row of Table 1.

For the frequentist analysis we achieve sparsity in the regression model using two methods; (i) standard stepwise regression (ii) L1 regularization penalty, Lasso. The results of these analyses appear in Table 2. The stepwise and Bayesian variable selection methods select identical variables for y =baseline NEET status. In stark contrast, the results from the Lasso analysis shows that the only variable with a p-value less than the traditional cut-off of 0.05, is sex.

The results of our partition method appear in Figures 1 and 2. Figures 1 and 2 show the posterior probabilities $\Pr(\gamma_{ks} = 1|y, w)$, with $w = (1, \text{Age}_s, G_s)$, where Age_s is age years, ranging from 16 to 25 and $G_s = 1$ corresponds to female and $G_s = 0$ corresponds to male.. The number of distinct control categories, $S = 20$. Figures 1 and 2 shows that the MPPs of the modifiable variables depends upon the control variables. In Figure 1, we see that the MPPs of alcohol, tobacco and depression all show a step-like increase from the ages of 19-21 onwards; that the MPPs of cannabis and tobacco spike at a few discrete ages; and that the MPP of functioning is high up to age 20 then has a step like decrease afterwards. In Figure 2, we see that the MPPs of alcohol, cannabis, tobacco and functioning all show MPP spikes at various discrete ages depending on sex; that depression again shows a step-like increase in MPP from age 20 onwards; and that anxiety shows a step-like decrease in MPP from age 20 onwards. The specific variables in the partition models are generally similar to those selected in the non-partition models which include the demographic variables. However, we note that a lot of the 'promising' variables are activated or have high MPPs at specific ages. The effects of age generally make intuitive sense, the MPPs of alcohol and tobacco having a more probable relationship with NEET as the age of participants increase. This is also confirmed by including interaction terms on the control covariates in the non-partition model, the results of which are consistent with the findings in the partition model.


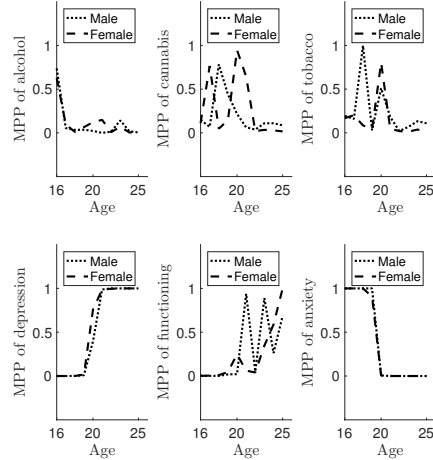
Our method is able to provide more specific inferences about important variables within a given partition and we see that our new formulation of priors provides meaningful differences in the interpretation of results. The results highlight the need to consider different intervention strategies for individuals with different non-modifiable factors.

Table 1: Bayesian variable selection with no covariate partition

Variable	(a) $y = \text{baseline NEET status}$				(b) $y = \text{follow up NEET status}$			
	(i) Ex. age, sex		(ii) Inc. age, sex		(i) Ex. age, sex		(ii) Inc. age, sex	
	$\hat{\beta}$	MPP	$\hat{\beta}$	MPP	$\hat{\beta}$	MPP	$\hat{\beta}$	MPP
Alcohol	-0.01	0.16	-0.03	0.97	0.00	0.08	0.00	0.17
Cannabis	0.02	0.98	0.02	0.97	0.01	0.57	0.01	0.48
Tobacco	0.00	0.02	0.00	0.02	0.00	0.03	0.00	0.02
Depression	0.06	1	0.06	1	0.02	0.55	0.03	0.53
Disability	0.00	0.09	0.00	0.12	0.01	0.52	0.01	0.57
Anxiety	0.00	0.03	0.00	0.02	0.00	0.02	0.00	0.02
Age			0.12	1			0.00	0.09
Sex			-0.61	1			-0.44	0.97
ROC AUC	0.64		0.72		0.58		0.60	

Table 2: Stepwise and Lasso results using $y = \text{baseline NEET status}$

Variable	(a) Stepwise results						(b) Lasso results					
	(i) Ex. age, sex			(ii) Inc. age, sex			(i) Ex. age, sex			(ii) Inc. age, sex		
	$\hat{\beta}$	SE	p	$\hat{\beta}$	SE	p	$\hat{\beta}$	SE	p	$\hat{\beta}$	SE	p
Alcohol				-0.03	0.01	0.01	-0.10	0.08	0.21	-0.27	0.15	0.08
Cannabis	0.02	0.01	0.02	0.02	0.01	0.02	0.13	0.08	0.07	0.22	0.14	0.11
Tobacco							0.03	0.16	0.85			
Depression	0.05	0.01	0.00	0.06	0.01	0.00	0.11	0.10	0.27	0.23	0.17	0.17
Functioning							0.12	0.10	0.22	0.23	0.17	0.18
Anxiety												
Age				0.12	0.03	0.00				0.16	0.14	0.22
Sex				-0.58	0.13	0.00				-0.32	0.13	0.01

Figure 1: Marginal posterior probabilities (MPP) over the different partitions of age and sex for each modifiable variable with $y = \text{baseline NEET status}$ Figure 2: Marginal posterior probabilities (MPP) over the different partitions of age and sex for each modifiable variable with $y = \text{followup NEET status}$

References

- Albert, J. and Chib, S. (1993), “Bayesian analysis of binary and polychotomous response data,” *Journal of the American Statistical Association*, 88, 669–679.
- Chernozhukov, V., Chetverikov, D., Demirer, M., Duflo, E., Hansen, C., and Newey, W. (2016), “Double machine learning for treatment and causal parameters,” *ArXiv e-prints*.
- Gelman, A. (2006), “Prior distributions for variance parameters in hierarchical models,” *Bayesian Analysis*, 1, 515–534.
- Lamnisos, D., Griffin, J., and Steel, M. (2009), “Transdimensional sampling algorithms for Bayesian variable selection in classification problems with many more variables than observations,” *Journal of Computational and Graphical Statistics*, 18, 592–612.
- O’Dea, B., Glozier, N., Purcell, R., McGorry, P., Scott, J., Feilds, K., Hermens, D., Buchanan, J., Scott, E., Yung, A., Killackey, E., Guastella, A., and Hickie, I. (2014), “A cross-sectional exploration of the clinical characteristics of disengaged (NEET) young people in primary mental healthcare,” *BMJ Open*, 4.
- Purcell, R., Jorm, A., Hickie, I., Yung, A., Pantelis, C., Amminger, G., Glozier, N., Killackey, E., Phillips, L., Wood, S., Mackinnon, A., Scott, E., Kenyon, A., Mundy, L., Nichles, A., Scaffidi, A., Spiliotacopoulos, D., Taylor, L., Tong, J., Wiltink, S., Zmicerevska, N., Hermens, D., Guastella, A., and McGorry, P. (2015), “Transitions Study of predictors of illness progression in young people with mental ill health: study methodology,” *Early Intervention in Psychiatry*, 9, 38–47.